UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/669,870	09/24/2003	Carlton A. Andrews	016295.1460	5137

23640 7590 12/14/2007

BAKER BOTTS, LLP
910 LOUISIANA
HOUSTON, TX 77002-4995

EXAMINER	
CHEA, PHILIP J	
ART UNIT	**PAPER NUMBER**
2153	

NOTIFICATION DATE	DELIVERY MODE
12/14/2007	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

debbie.allen@bakerbotts.com

Office Action Summary	Application No.	Applicant(s)	
	10/669,870	ANDREWS ET AL.	
	Examiner	Art Unit	
	Philip J. Chea	2153	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *20 September 2007*.

2a)☒ This action is **FINAL**. 2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-15,17-49 and 51-59* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-15,17-49 and 51-59* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☐ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date _____.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____.
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

This Office Action is in response to an Amendment filed September 20, 2007. Claims 1-15,17-49,51-59 are currently pending. Any rejection not set forth below has been overcome by the current Amendment.

Claim Rejections - 35 USC § 103

1. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

2. Claims 1-2, and 21 are rejected under 35 USC 103 as being unpatentable over Scharber (Pat # US 6,542,964) in view of Squire et al (Pub # 2002/0049840), and further in view of Kishi (US 6,163,773).

In reference to claims 1-2, and 21 Scharber teaches a limitation for a server having one or more systems capable of implementing two or more cache policies (see spec, sec. 6, lines 21-27, which implies this limitation because the invention teaches a cache server implemented for the selection of one of a plurality of cache protocols), a network operative with said server (see spec, sec. 5, lines 53-60, which teaches this limitation because the cache protocol is selected for traffic conditions in a network), said network connecting one or more clients to said server (see spec, sec. 7, lines 1-7, which teaches this limitation because the cache server can accessed across the network by one or more clients), said clients constructed and arranged to communicate with said server thereby placing a load on said server (see spec, sec. 5, lines 25-30, which implies this limitation because client requests to the same server cause the server's load to increase), a load monitor constructed and arranged to monitor said load (see spec, sec. 8, lines 56-65, which implies this limitation because the load balancing procedure within the invention is based on monitoring IP statistics, such as the overall weight and load of each server), said load monitor further constructed and arranged to select one or more cache policies that optimize a performance characteristic of said information handling system (see spec, sec. 5, lines 53-60, which implies this

limitation because a cache protocol is selected based on load balancing requirements within the network), wherein said load manager monitors said load and implements a cache policy that optimizes a characteristic of said information handling system (see spec, sec. 6, lines 17-21, which teaches this limitation because the cache protocol selected by the load balancer is used to mange stored content on an internet content delivery system),wherein said load monitor employs a template to select one or more cache policies (see spec, sec. 5, lines 53-60, which implies this limitation because a cache protocol is selected using a schema including type of content, site associated with the content, server resource ability, and class of service requirements), and wherein said load monitor is a load balancer (see spec, sec. 8, lines 60-65, which teaches this limitation because a load balancer is used to examine the weights of each server in order to choose the correct cache protocol).

Scharber teaches all the limitations as disclosed above except for a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems.

The general concept of a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems is well known in the art as illustrated by Squire et al, which teach the limitation for a load monitor operative with one or more Raid systems (see e.g. [0037], lines 1-9 & [0038], lines 1-10, which implies this limitation because the network cache of the load balancing device used to connect clients to the server (as shown in sec. 0011, lines 6-13) may be implemented as a Raid system).

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber to include the use of a limitation for a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems as illustrated by Squire et al in order to effectively implement a network caching system, as implied in sec. [0038], lines 12-14 of Squire et al.

Although the system disclosed by Scharber in view of Squire shows substantial features of the claimed invention (discussed above), it fails to disclose wherein at least one of said cache policies is an adaptive policy based on previous activity in the information handling system.

Nonetheless, these features are well known in the art and would have been an obvious modification of the system disclosed by Scharber in view of Squire, as evidenced by Kishi.

In an analogous art, Kishi discloses an information handling system (i.e. data storage system) where a cache is managed by a predictive cache management engine (see Abstract). Further disclosing a cache policy that is an adaptive policy based on previous activity in the information handling system (see column 5, lines 4-9, describing a training (i.e. adaptive) to rank cached datasets based upon past activity).

Given the teaching of Kishi, a person having ordinary skill in the art would have readily recognized the desirability and advantages of modifying Scharber in view of Squire by employing an adaptive cache policy, such as disclosed by Kishi, in order to determine when data should be cached or not (see Kishi column 1, lines 56-58).

3. Claims 3-5, 8-9, 17, and 20 are rejected under 35 USC 103 as being unpatentable over Scharber (Pat # US 6,542,964) and Squire et al (Pub # 2002/0049840) and Kishi as applied to claim 1 and further in view of Patel et al (Pat # US 7,146,524).

In reference to claims 3-5, 8-9, 17, and 20 Scharber teaches a limitation for a load monitor constructed and arranged to monitor said load (see spec, sec. 8,as stated above).
Scharber and Squire et al teach all the limitations as disclosed above except for a load monitor which employs an algorithm to select said one or more cache policies, said load monitor employing a template and an algorithm to select said one or more cache policies, a RAID system having a read cache, said read cache having a read-ahead policy, said RAID system having a write cache, wherein one of said cache policies is a read-ahead policy, and wherein a cache policy is a cached policy.

The general concept of providing a limitation of a load monitor employing an algorithm to select said one or more cache policies, a load monitor employing a template and an algorithm to select said one or more cache policies, a RAID system having a read cache, a read cache having a read-ahead policy, a RAID system having a write cache, wherein one of said cache policies is a read-ahead policy, and wherein said cache policy is a cached policy are well known in the art as illustrated by Patel et al, which

teach the limitation for wherein said load monitor employs an algorithm to select said one or more cache policies (see spec, sec. 14, lines 35-42, which teaches this limitation because an algorithm is used to determine the appropriate caching scheme to be implemented within the load balancing system), wherein said load monitor employs a template and an algorithm to select said one or more cache policies (see spec, sec. 14, lines 35-42, which teaches this limitation because an algorithm is used to determine the appropriate caching scheme to be implemented and sec. 15, lines 5-14, which teaches this limitation because the modifications made to cache protocols are used based on the type of data being stored, the processing speed, and the number of storage units within the load balance switching system), wherein said RAID system has a read cache (see spec, sec. 2, lines 4-6 and sec. 14, lines 59-61, which teaches this limitation because device within the Raid system may read data and the cache module may perform read aheads), wherein said read cache has a read-ahead policy (see spec, sec. 14, lines 59-61, which teaches this limitation because the cache module may perform read aheads), wherein said RAID system has a write cache (see spec, sec. 14, lines 49-53, which teaches this limitation because the caching schemes used allow for disk write implementations), wherein one of said cache policies is a read-ahead policy (see spec, sec. 14, lines 59-61, which teaches this limitation because the cache module may perform read aheads), and wherein said cache policy is a cached policy (see spec, sec. 14, lines 53-55, which teaches this limitation because different caching schemes, such as the last recently used scheme of sec. 14, lines 43-45, may implement different caching protocols).

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber and Squire et al to include the use of a limitation for wherein said load monitor employs an algorithm to select said one or more cache policies, wherein said load monitor employs a template and an algorithm to select said one or more cache policies, wherein said RAID system has a read cache, wherein said read cache has a read-ahead policy, wherein said RAID system has a write cache, wherein one of said cache policies is a read-ahead policy, and wherein said cache policy is a cached policy as illustrated by Patel et al in order to improve upon implementing a variety of caching protocols, as implied in sec. 15, lines 1-4 of Patel et al.

4. Claims 6-7, 10-15, and 19-20 are rejected under 35 USC 103 as being unpatentable over Scharber (Pat # US 6,542,964) and Squire et al (Pub # 2002/0049840) and Patel et al (Pat # US 7,146,524) in view of Kishi as applied to claim 5 and further in view of Surugucchi et al (Pub # US 2002/0095532).

In reference to claims 6-7, 10-16, and 19-20 Scharber teaches a limitation for a load monitor constructed and arranged to monitor said load (see spec, sec. 8, as stated above).
Scharber, Squire et al, and Patel et al teach all the limitations as disclosed above except for a cache policy being a no-ahead policy, wherein one of said cache policies is an adaptive policy, wherein one of said cache policies is back policy, and wherein one of said cache policy is a through policy.

The general concept of providing a limitation of a read cache having a no-ahead policy, wherein said read cache has an adaptive policy, wherein said write cache has back policy, wherein said write cache has through policy, wherein said RAID system has an I/O, wherein said I/O has a cached policy, wherein said I/O has a direct policy, wherein one of said cache policies is a no-ahead policy, and wherein one of said cache policies is back policy, and wherein one of said cache policy is a through policy is well known in the art as illustrated by Surugucchi et al, which teach a limitation of a read cache having a no-ahead policy (see e.g. [0009], lines 9-13, which teaches this limitation because the invention has a parameter to determine whether or not the read-ahead caching function is implemented or disabled and the only positive element that the applicant gave for a no-ahead policy was that the disk controller didn't implement a read-ahead policy), wherein said write cache has a back policy (see e.g. [0009], lines 11-13, which implies this limitation because a write-back caching policy is implemented), wherein said write cache has through policy (see e.g. [0009], lines 11-13, which implies this limitation because write-through caching is implemented within the system that utilizes Raid controllers), wherein said RAID system has an I/O (see e.g. [0010], lines 1-3, which teaches this limitation because cached i/o is embedded within the raid controller system), wherein said I/O has a cached policy (see e.g. [0010], line 3, which teaches this limitation because cached i/o is implemented), wherein said I/O has a direct policy (see e.g. [0010], line 3, which teaches this limitation because direct i/o is implemented), wherein one of said cache policies is a no-ahead policy (see e.g. [0009], line 11, which teaches this limitation because the invention has a

parameter to determine whether or not the read-ahead caching function is implemented or disabled and the only positive element that the applicant gave for a no-ahead policy was that the disk controller didn't implement a read-ahead policy), wherein one of said cache policies is an adaptive policy (see e.g. [0004], which implies this limitation because the only positive element that the applicant gave of and 'adaptive' read cache policy was that it is typically the default setting for a read cache and the prior reads on this claim because is shows in lines 22-23 that the retrieval of data from cache are usually done via fast electronic RAM), wherein one of said cache policies is back policy (see e.g. [0009], line 12, which implies this limitation because a write-back caching policy is implemented), and wherein one of said cache policy is a through policy (see e.g. [0009], line 12, which implies this limitation because write-through caching is implemented within the system that utilizes Raid controllers).

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber, Squire et al, and Patel et al to include the use of a limitation for wherein said read cache has a no-ahead policy, wherein said read cache has an adaptive policy, wherein said write cache has back policy, wherein said write cache has through policy, wherein said RAID system has an I/O, wherein said I/O has a cached policy, wherein said I/O has a direct policy, wherein one of said cache policies is a no-ahead policy, wherein one of said cache policies is an adaptive policy, and wherein one of said cache policies is back policy, and wherein one of said cache policy is a through policy are well known in the art as illustrated by Surugucchi I et al in order to effectively implement the use of a raid controller, as implied in sec. [0009], lines 1-4 of Surugucchi et al.

5. Claims 22, 23 and 24 are rejected under 35 USC 103 as being unpatentable over Scharber (Pat # US 6,542,964) in view of Squire et al (Pub # 2002/0049840) in view of Kishi.

In reference to claims 22, 23 and 24 Scharber teaches a limitation for a load monitor constructed and arranged to monitor said load (see spec, sec. 8, as specified above)

Scharber teaches all the limitations as disclosed above except for a server having one or more RAID systems and wherein the load monitor is a router, server, or a cluster master.

The general concept of a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems is well known in the art as illustrated by Squire et al, which teach the limitation for a load monitor operative with one or more Raid systems (see e.g. [0037], as stated above).

The general concept of a limitation of wherein a load monitor is a router, server, or a cluster master as opposed to a load balancer is rejected under obvious design optimization because one of ordinary skill in the art would find it obvious to implement any of the three embodiments in place of a load balancer since a load balancer acts a cluster manager by evenly distributing the load or weight of client connections throughout a cluster of servers, a load balancer maintains even weight amongst servers, therefore using a server as a load monitor would be perceived to be obvious, and it would be obvious to implement a router as a load monitor in place of a load balancer as a load balancer also evenly distributes client requests to servers transmitted through a router.

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber to include the use of a limitation for a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems as illustrated by Squire et al in order to effectively implement a network caching system, as implied in sec. [0038], lines 12-14 of Squire et al.

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber to include the use of a limitation for wherein the load monitor is a router, server, or a cluster master under the scope of obvious design optimization in order to select a cache protocol based on load balancing requirements and traffic conditions in a network, as implied in see spec, sec. 5, lines 53-60 of Scharber.

6. Claims 25 and 26 are rejected under 35 USC 103 as being unpatentable over Scharber (Pat # US 6,542,964) in view of Squire et al (Pub # 2002/0049840) in view of Kishi.

In reference to claims 25 and 26 Scharber teaches a limitation for at least one server having one or more systems capable of implementing two or more cache policies (see spec, sec. 6, lines 21-27,

which implies this limitation because the invention teaches a cache server implemented for the selection of one of a plurality of cache protocols), a network operative with said server (see spec, sec. 5, lines 53-60, which teaches this limitation because the cache protocol is selected for traffic conditions in a network), said network connecting one or more clients to said server (see spec, sec. 7, lines 1-7, which teaches this limitation because the cache server can accessed across the network by one or more clients), said clients constructed and arranged to communicate with said server thereby placing a load on said server (see spec, sec. 5, lines 25-30, which implies this limitation because client requests to the same server cause the server's load to increase), a load monitor constructed and arranged to monitor said load (see spec, sec. 8, lines 56-65, which implies this limitation because the load balancing procedure within the invention is based on monitoring IP statistics, such as the overall weight and load of each server), said load monitor further constructed and arranged to select one or more cache policies that optimize a performance characteristic of said information handling system (see spec, sec. 5, lines 53-60, which implies this limitation because a cache protocol is selected based on load balancing requirements within the network), wherein said load manager monitors said load and implements a cache policy that optimizes a characteristic of said information handling system (see spec, sec. 6, lines 17-21, which teaches this limitation because the cache protocol selected by the load balancer is used to mange stored content on an internet content delivery system), and wherein said load monitor is a load balancer (see spec, sec. 8, lines 60-65, which teaches this limitation because a load balancer is used to examine the weights of each server in order to choose the correct cache protocol).

Scharber teaches all the limitations as disclosed above except for a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems.

The general concept of a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems is well known in the art as illustrated by Squire et al, which teach the limitation for a load monitor operative with one or more Raid systems (see e.g. [0037], lines 1-9 & [0038], lines 1-10, which implies this limitation because the network

cache of the load balancing device used to connect clients to the server (as shown in sec. 0011, lines 6-13) may be implemented as a Raid system).

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber to include the use of a limitation for a server having one or more RAID systems, a load monitor operative with one or more Raid systems, and cache policies of one or more raid systems as illustrated by Squire et al in order to effectively implement a network caching system, as implied in sec. [0038], lines 12-14 of Squire et al.

Although the system disclosed by Scharber in view of Squire shows substantial features of the claimed invention (discussed above), it fails to disclose wherein at least one of said cache policies is an adaptive policy based on previous activity in the information handling system.

Nonetheless, these features are well known in the art and would have been an obvious modification of the system disclosed by Scharber in view of Squire, as evidenced by Kishi.

In an analogous art, Kishi discloses an information handling system (i.e. data storage system) where a cache is managed by a predictive cache management engine (see Abstract). Further disclosing a cache policy that is an adaptive policy based on previous activity in the information handling system (see column 5, lines 4-9, describing a training (i.e. adaptive) to rank cached datasets based upon past activity).

Given the teaching of Kishi, a person having ordinary skill in the art would have readily recognized the desirability and advantages of modifying Scharber in view of Squire by employing an adaptive cache policy, such as disclosed by Kishi, in order to determine when data should be cached or not (see Kishi column 1, lines 56-58).

7. Claims 27, 28 and 29 are rejected under 35 USC 103 as being unpatentable over Scharber (Pat # US 6,542,964) in view of Squire et al (Pub # 2002/0049840), in view of Kishi.

In reference to claims 27, 28 and 29 Scharber teaches a limitation for a load monitor constructed and arranged to monitor said load (see spec, sec. 8, as specified above).

Scharber and Squire et al teach all the limitations as disclosed above except for a server having one or more RAID systems and wherein the load monitor is a router, server, or a cluster master.

The general concept of a limitation for wherein the load monitor is a router, server, or a cluster master as opposed to a load balancer is rejected under obvious design optimization because one of ordinary skill in the art would find it obvious to implement any of the three embodiments in place of a load balancer since a load balancer acts a cluster manager by evenly distributing the load or weight of client connections throughout a cluster of servers, a load balancer maintains even weight amongst servers, therefore using a server as a load monitor would be perceived to be obvious, and it would be obvious to implement a router as a load monitor in place of a load balancer as a load balancer also evenly distributes client requests to servers transmitted through a router.

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Scharber and Squire et al to include the use of a limitation for wherein the load monitor is a router, server, or a cluster master under the scope of obvious design optimization in order to select a cache protocol based on load balancing requirements and traffic conditions in a network, as implied in see spec, sec. 5, lines 53-60 of Scharber.

8. Claims 30-33, 35-39, 42-43, 51, and 54-55 are rejected under 35 U.S.C. 103(a) as being unpatentable over Patel et al (Pat # US 7,146,524), and further in view of Kishi (US 6,163,773).

Patel et al teach a method for changing cache policy for a raid system on an information system (see spec, sec. 15, lines 1-4, which teaches this limitation because the caching protocol utilized within this filing system may be varied depending on system requirements, note that the filing system may be implemented using a raid storage system as shown in sec. 1, lines 57-59), reading a set of templates (see spec, sec. 15, lines 5-14, which teaches this limitation because the modifications made to cache protocols are used based on the type of data being stored, the processing speed, and the number of storage units), determining the load of the network (see spec, sec. 13, lines 22-25, which teaches this limitation because load balancing techniques are used to balance the data blocks of the storage units within the network, also note sec. 21, lines 40-47 which shows how total storage capacity for each data

block of the network is tracked), indexing said templates using said load to determine a cache setting (see spec, sec. 18, lines 17-22, which teaches this limitation because the data blocks being stored are indexed into 4 type data locations, shown in sec. 18, lines 1-5, based on the caching protocol requirements, shown in sec. 15, lines 5-14, that each data block must be stored according to), and applying said cache settings to said RAID system (see spec, sec. 16, lines 20-25, which teaches this limitation because the selected cache is used to improve the storage unit's (which may be a raid storage device, as shown in sec. 1, lines 55-60) performance in retrieving data blocks).

Although the system disclosed by Patel shows substantial features of the claimed invention (discussed above), it fails to disclose wherein at least one of said cache policies is an adaptive policy based on previous activity in the information handling system.

Nonetheless, these features are well known in the art and would have been an obvious modification of the system disclosed by Patel, as evidenced by Kishi.

In an analogous art, Kishi discloses an information handling system (i.e. data storage system) where a cache is managed by a predictive cache management engine (see Abstract). Further disclosing a cache policy that is an adaptive policy based on previous activity in the information handling system (see column 5, lines 4-9, describing a training (i.e. adaptive) to rank cached datasets based upon past activity).

Given the teaching of Kishi, a person having ordinary skill in the art would have readily recognized the desirability and advantages of modifying Patel by employing an adaptive cache policy, such as disclosed by Kishi, in order to determine when data should be cached or not (see Kishi column 1, lines 56-58).

With respect to claim 31, Patel et al teach a method for invoking a delay (see spec, sec. 5, lines 54-58, which teaches this limitation because access to information within the storage units may be delayed until the data can be restored from backup tapes).

With respect to claim 32, Patel et al teach a method for wherein steps are performed continuously on said information handling system (see spec, sec. 6, lines 30-32, which teaches this limitation because the file system may be set to automatically complete tasks using the storage units).

With respect to claim 33, Patel et al teach a method for wherein steps are performed continuously in real time on said information handling system (see spec, sec. 2, lines 15-20, which teaches this limitation because data replication (caching) and movement embodied within the invention are performed in real time).

With respect to claim 35, Patel et al teach a method for wherein said step of determining employs a template to select said cache setting (see spec, sec. 15, lines 5-14, which teaches this limitation because the modifications made to cache protocols are used based on the type of data being stored, the processing speed, and the number of storage units).

With respect to claim 36, Patel et al teach a method for wherein said step of determining employs an algorithm to select said cache setting (see spec, sec. 14, lines 35-42, which teaches this limitation because an algorithm is used to determine the appropriate caching scheme to be implemented).

With respect to claim 37, Patel et al teach a method for wherein said step of determining employs a template and an algorithm to select said cache setting (see spec, sec. 14, lines 35-42, which teaches this limitation because an algorithm is used to determine the appropriate caching scheme to be implemented and sec. 15, lines 5-14, which teaches this limitation because the modifications made to cache protocols are used based on the type of data being stored, the processing speed, and the number of storage units).

With respect to claim 38, Patel et al teach a method for wherein said cache setting is made up of two or more policies (see spec, sec. 14, lines 54-56, which teaches this limitation because different cache protocols such as on demand or read ahead may be implemented).

With respect to claim 39, Patel et al teach a method for wherein said RAID system has a read cache (see spec, sec. 2, lines 4-6 and sec. 14, lines 59-61, which teaches this limitation because device within the Raid system may read data and the cache module may perform read aheads).

With respect to claim 42 and 51, Patel et al teach a method for wherein said read cache and cache setting has a read-ahead policy (see spec, sec. 14, lines 59-61, which teaches this limitation because the cache module may perform read aheads).

With respect to claim 43, Patel et al teach a method for wherein said RAID system has a write cache (see spec, sec. 14, lines 49-53, which teaches this limitation because the caching schemes used allow for disk write implementations).

With respect to claim 54, Patel et al teach a method for wherein said cache setting includes a cached policy (see spec, sec. 14, lines 53-55, which teaches this limitation because different caching schemes, such as the last recently used scheme of sec. 14, lines 43-45, may implement different caching protocols).

With respect to claim 55, Patel et al teach a method for wherein said load monitor is a load balancer (see spec, sec. 7, lines 29-35, which teaches this limitation because a load balancing switch is used to balance the data block requests directed to an application server).

9. Claims 40-41, 44-49 and 52-53 are rejected under 35 USC 103 as being unpatentable over Patel et al (Pat # US 7,146,524) in view of Surugucchi et al (Pub # US 2002/0095532).

In reference to claims 40-41, 44-49, and 52-53 Patel et al teach a limitation for wherein said RAID system has a read cache (see spec, sec. 2, as stated above).

Patel et al teach all the limitations as disclosed above except for a read cache having a no-ahead policy, wherein said read cache has an adaptive policy, wherein said write cache has back policy, wherein said write cache has through policy, wherein said RAID system has an I/O, wherein said I/O has a cached policy, wherein said I/O has a direct policy, wherein one of said cache policies is a no-ahead policy, wherein one of said cache policies is an adaptive policy, and wherein one of said cache policies is back policy, and wherein one of said cache policy is a through policy.

The general concept of providing a limitation of a read cache having a no-ahead policy, wherein said read cache has an adaptive policy, wherein said write cache has back policy, wherein said write cache has through policy, wherein said RAID system has an I/O, wherein said I/O has a cached policy, wherein said I/O has a direct policy, wherein one of said cache policies is a no-ahead policy, and wherein one of said cache policies is back policy, and wherein one of said cache policy is a through policy is well known in the art as illustrated by Surugucchi et al, which teach a limitation for wherein said read cache

has a no-ahead policy (see e.g. [0009], lines 9-13, which teaches this limitation because the invention has a parameter to determine whether or not the read-ahead caching function is implemented or disabled and the only positive element that the applicant gave for a no-ahead policy was that the disk controller didn't implement a read-ahead policy) and wherein said write cache has a back policy (see e.g. [0009], lines 11-13, which implies this limitation because a write-back caching policy is implemented), wherein said write cache has through policy (see e.g. [0009], lines 11-13, which implies this limitation because write-through caching is implemented within the system that utilizes Raid controllers), wherein said RAID system has an I/O (see e.g. [0010], lines 1-3, which teaches this limitation because cached i/o is embedded within the raid controller system), wherein said I/O has a cached policy (see e.g. [0010], line 3, which teaches this limitation because cached i/o is implemented), wherein said I/O has a direct policy (see e.g. [0010], line 3, which teaches this limitation because direct i/o is implemented), wherein one of said cache policies is a no-ahead policy (see e.g. [0009], line 11, which teaches this limitation because the invention has a parameter to determine whether or not the read-ahead caching function is implemented or disabled and the only positive element that the applicant gave for a no-ahead policy was that the disk controller didn't implement a read-ahead policy), wherein one of said cache policies is an adaptive policy (see e.g. [0004], which implies this limitation because the only positive element that the applicant gave of and 'adaptive' read cache policy was that it is typically the default setting for a read cache and the prior reads on this claim because is shows in lines 22-23 that the retrieval of data from cache are usually done via fast electronic RAM), wherein one of said cache policies is back policy (see e.g. [0009], line 12, which implies this limitation because a write-back caching policy is implemented), and wherein one of said cache policy is a through policy (see e.g. [0009], line 12, which implies this limitation because write-through caching is implemented within the system that utilizes Raid controllers).

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Patel et al to include the use of a limitation for a read cache having a no-ahead policy, wherein said read cache has an adaptive policy, wherein said write cache has back policy, wherein said write cache has through policy, wherein said RAID system has an I/O, wherein said I/O has a cached policy, wherein said I/O has a direct policy, wherein one of said cache policies is a no-ahead policy, wherein one of said cache

policies is an adaptive policy, and wherein one of said cache policies is back policy, and wherein one of said cache policy is a through policy are well known in the art as illustrated by Surugucchi I et al in order to effectively implement the use of a raid controller, as implied in sec. [0009], lines 1-4 of Surugucchi et al.

10. Claims 34, and 56-58 are rejected under 35 USC 103 as being unpatentable over Patel et al (Pat # US 7,146,524) as applied to claim 30 and further in view of obvious design optimization.

In reference to claims 56, 57 and 58 Patel et al teach a limitation for wherein the load monitor is a load balancer (see spec, sec. 7, as stated above) and wherein steps are performed continuously in real time on said information handling system (see spec, sec. 2, as stated above).

Patel et al teach all the limitations as disclosed above except for wherein the load monitor is a router, server, or a cluster master and wherein the steps are performed continuously in near real time on said information handling system.

The general concept of a limitation for wherein the load monitor is a router, server, or a cluster master as opposed to a load balancer is rejected under obvious design optimization because one of ordinary skill in the art would find it obvious to implement any of the three embodiments in place of a load balancer since a load balancer acts a cluster manager by evenly distributing the load or weight of client connections throughout a cluster of servers, a load balancer maintains even weight amongst servers, therefore using a server as a load monitor would be perceived to be obvious, and it would be obvious to implement a router as a load monitor in place of a load balancer as a load balancer also evenly distributes client requests to servers transmitted through a router. One of ordinary skill in the art would find it obvious to implement a limitation for wherein the steps are performed continuously in near real time on said information handling system because Patel et al specifies in sec. 2, lines 15-20 that the data replication aspects for a caching system are performed in real-time. It would have been obvious for one to implement the system in a near-real-time aspect rather than a strictly real-time embodiment.

It would have been obvious for one of ordinary skill in the art at the time of the invention to modify Patel et al to include the use of a limitation for wherein the load monitor is a router, server, or a cluster master and wherein the steps are performed continuously in near real time on said information handling

system under the scope of obvious design optimization in order to successfully implement a load balance switching system, as implied in sec. 7, lines 29-35 of Patel et al.

11. Claim 59 rejected under 35 U.S.C. 103(a) as being unpatentable over Grosner et al (Pat # US 7,089,293), and further in view of Kishi (US 6,163,773).

Grosner et al teach a method including a limitation for providing client means for generating a load (see spec, sec. 21, lines 1-3, which teaches this limitation because the invention has a mechanism to balance requests, which generate a load, amongst servers in the system), server means for servicing said load (see spec, sec. 21, lines 1-3, which teaches this limitation because the invention has a mechanism to balance requests amongst server), said sever having one or more RAID systems (see spec, sec. 24, lines 43-47, which teaches this limitation because servers handle different Raid replications of the same file system), said RAID systems capable of implementing two or more cache policies (see spec, sec. 20, lines 34-36, which teaches this limitation because the raid file system replication system is capable of utilizing one or more disk caching modules, such as object-based caching), network means operative with said server means and said client means for transmitting said load (see spec, sec. 24, lines 43-47 which teaches this limitation because the load is transmitted in workload distribution fashion amongst a plurality of server), load monitor means for monitoring said load (see spec, sec. 24, lines 43-47, which teaches this limitation because a load balancing mechanism is implemented to even distribute workload throughout the plurality of server), said load monitor means operative with said one or more RAID systems (see spec, sec. 24, lines 43-47, which teaches this limitation because the load balancing mechanism handles Raid replications of file systems), said load monitor constructed and arranged to select a cache policy of said one or more RAID systems that optimize a performance characteristic of said information handling system (see spec, sec. 27, 9-11, which teaches this limitation because a lb (load balancer) can shift between implementing maximally distributed caching techniques or other performance enhancing techniques within the data caching system), and wherein said load manager monitors said load and implements a cache policy that optimizes a characteristic of said information handling system (see spec, sec. 20, lines 34-38, which teaches this limitation because different caching techniques may

be implemented within the load balancing system and algorithms may be implemented to replace the current disk cache model).

Although the system disclosed by Grosner shows substantial features of the claimed invention (discussed above), it fails to disclose wherein at least one of said cache policies is an adaptive policy based on previous activity in the information handling system.

Nonetheless, these features are well known in the art and would have been an obvious modification of the system disclosed by Grosner, as evidenced by Kishi.

In an analogous art, Kishi discloses an information handling system (i.e. data storage system) where a cache is managed by a predictive cache management engine (see Abstract). Further disclosing a cache policy that is an adaptive policy based on previous activity in the information handling system (see column 5, lines 4-9, describing a training (i.e. adaptive) to rank cached datasets based upon past activity).

Given the teaching of Kishi, a person having ordinary skill in the art would have readily recognized the desirability and advantages of modifying Grosner by employing an adaptive cache policy, such as disclosed by Kishi, in order to determine when data should be cached or not (see Kishi column 1, lines 56-58).

Response to Arguments

12. Applicant's arguments with respect to claims 1-15,17-49,51-59 have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

13. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date

of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Philip J. Chea whose telephone number is 571-272-3951. The examiner can normally be reached on M-F 6:30-4:00 (1st Friday Off).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Glenn Burgess can be reached on 571-272-3949. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Philip J Chea
Examiner
Art Unit 2153

PJC 11/27/07

GLENTON B. BURGESS
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100